Strategy Improvement and Randomized Subexponential Algorithms for Stochastic Parity Games

Krishnendu Chatterjee[1] and Thomas A. Henzinger[1,2]

[1] EECS, UC Berkeley, USA
[2] EPFL, Switzerland
{c_krish,tah}@eecs.berkeley.edu

Abstract. A stochastic graph game is played by two players on a game graph with probabilistic transitions. We consider stochastic graph games with ω-regular winning conditions specified as parity objectives. These games lie in NP $\cap$ coNP. We present a strategy improvement algorithm for stochastic parity games; this is the first non-brute-force algorithm for solving these games. From the strategy improvement algorithm we obtain a randomized subexponential-time algorithm to solve such games.

1 Introduction

Graph games. A stochastic graph game [5] is played on a directed graph with three kinds of states: player-1, player-2, and probabilistic states. At player-1 states, player 1 chooses a successor state; at player-2 states, player 2 chooses a successor state; at probabilistic states, a successor state is chosen according to a given probability distribution. The outcome of playing the game forever is an infinite path through the graph. If there are no probabilistic states, we refer to the game as a *2-player graph game*; otherwise, as a *2$^1\!/_2$-player graph game.*

Parity objectives. The theory of graph games with ω-regular winning conditions is the foundation for modeling and synthesizing reactive processes with fairness constraints. In the case of 2$^1\!/_2$-player graph games, the two players represent a reactive system and its environment, and the probabilistic states represent uncertainty. The *parity* objectives provide an adequate model, as the fairness constraints of reactive processes are ω-regular, and every ω-regular winning condition can be specified as a parity objective [11]. The solution problem for a 2$^1\!/_2$-player game with parity objective Φ asks for each state s, for the maximal probability with which player 1 can ensure the satisfaction of Φ if the game is started from s (this probability is called the *value* of the game at s). An *optimal strategy* for player 1 is a strategy that enables player 1 to win with that maximal probability. The existence of *pure memoryless* optimal strategies for 2$^1\!/_2$-player games with parity objectives was established recently in [4] (a pure memoryless strategy chooses for each player-1 state a unique successor state). The existence of pure memoryless optimal strategies implies that the solution problem for 2$^1\!/_2$-player games with parity objectives lies in NP $\cap$ coNP.

Previous algorithms. Emerson and Jutla [7] had showed in 1988 that 2-player parity games (*without* probabilistic states) can be solved in NP $\cap$ coNP. However, to date no polynomial-time algorithm is known to solve these games. In 2000, Vöge and Jurdziński [12] gave a *strategy improvement* algorithm for 2-player parity games. A strategy improvement scheme iterates local optimizations of a pure memoryless strategy; this works if the iteration can be shown to converge to a global optimum [8]. Although the best known bound for the worst-case running time of Vöge and Jurdziński is exponential, it behaves very well in practice. Moreover, Björklund et al. [1] used the strategy improvement scheme to derive a randomized subexponential-time algorithm for 2-player parity games. And recently, Jurdziński et al. [9] found a deterministic subexponential-time algorithm for 2-player parity games.

For $2^1\!/_2$-player games (*with* probabilistic states), Condon [5] proved containment in NP $\cap$ coNP in 1992 for the restricted case of *reachability* objectives, and she gave a strategy improvement algorithm for this subclass of $2^1\!/_2$-player games. Again, no polynomial-time algorithm is known to solve these games, but using strategy improvement, Ludwig [10] derived a randomized subexponential-time algorithm for $2^1\!/_2$-player reachability games on *binary* game graphs (game graphs with maximum out-degree 2). The techniques of [1] also yield a randomized subexponential-time algorithm for the nonbinary class of $2^1\!/_2$-player reachability games. However, the techniques of [9] do not extend to give a deterministic subexponential-time algorithm for $2^1\!/_2$-player reachability games. For the full class of $2^1\!/_2$-player games with general parity objectives, no algorithm has been known which is better than a brute-force enumeration of the set of all possible pure memoryless strategies, and chosing the best one.

Our results. We present the first strategy improvement algorithm for $2^1\!/_2$-player parity games. Our algorithm combines both techniques for 2-player parity games and for $2^1\!/_2$-player reachability games, employing a novel reduction from $2^1\!/_2$-player parity games (with quantitative winning criteria) to 2-player parity games (with qualitative winning criteria). We then show how the techniques of [1] can be extended to our strategy improvement algorithm to obtain a randomized subexponential algorithm for $2^1\!/_2$-player parity games. Given a game graph with n states and a parity objective with d priorities, the expected running time of our algorithm is $2^{O\left(\sqrt{d \cdot n \cdot \log(n)}\right)}$. The algorithm is subexponential if $d = O\left(\frac{n^{1-\varepsilon}}{\log(n)}\right)$ for some $\varepsilon > 0$. Thus, for the special case of reachability objectives, the expected running time of our algorithm matches the bound of the best known algorithm.

2 Definitions

We consider turn-based probabilistic games and some of its subclasses.

Game graphs. A *turn-based probabilistic game graph* (*$2^1\!/_2$-player game graph*) $G = ((S, E), (S_1, S_2, S_\bigcirc), \delta)$ consists of a directed graph (S, E), a partition $(S_1, S_2, S_\bigcirc)$ of the finite set S of states, and a probabilistic transition function δ: $S_\bigcirc \to \mathcal{D}(S)$, where $\mathcal{D}(S)$ denotes the set of probability distributions over the state space S. The states in S_1 are the *player*-1 states, where player 1 decides the successor state; the states in S_2 are the *player*-2 states, where player 2 decides the successor state; and the states in $S_{\bigcirc}$ are the *probabilistic* states, where the successor state is chosen according to the probabilistic transition function δ. We assume that for $s \in S_{\bigcirc}$ and $t \in S$, we have $(s,t) \in E$ iff $\delta(s)(t) > 0$, and we often write $\delta(s,t)$ for $\delta(s)(t)$. For technical convenience we assume that every state in the graph (S, E) has at least one outgoing edge. For a state $s \in S$, we write $E(s)$ to denote the set $\{t \in S \mid (s,t) \in E\}$ of possible successors. The *turn-based deterministic game graphs* (*2-player game graphs*) are the special case of the $2^1/_2$-player game graphs with $S_{\bigcirc} = \emptyset$. The *Markov decision processes* ($1^1/_2$-*player game graphs*) are the special case of the $2^1/_2$-player game graphs with $S_1 = \emptyset$ or $S_2 = \emptyset$. We refer to the MDPs with $S_2 = \emptyset$ as *player*-1 MDPs, and to the MDPs with $S_1 = \emptyset$ as *player*-2 MDPs.

Plays and strategies. An infinite path, or a *play*, of the game graph G is an infinite sequence $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ of states such that $(s_k, s_{k+1}) \in E$ for all $k \in \mathbb{N}$. We write Ω for the set of all plays, and for a state $s \in S$, we write $\Omega_s \subseteq \Omega$ for the set of plays that start from the state s. A *strategy* for player 1 is a function $\sigma\colon S^* \cdot S_1 \to \mathcal{D}(S)$ that assigns a probability distribution to all finite sequences $\boldsymbol{w} \in S^* \cdot S_1$ of states ending in a player-1 state (the sequence represents a prefix of a play). Player 1 follows the strategy σ if in each player-1 move, given that the current history of the game is $\boldsymbol{w} \in S^* \cdot S_1$, she chooses the next state according to the probability distribution $\sigma(\boldsymbol{w})$. A strategy must prescribe only available moves, i.e., for all $\boldsymbol{w} \in S^*$, $s \in S_1$, and $t \in S$, if $\sigma(\boldsymbol{w} \cdot s)(t) > 0$, then $(s,t) \in E$. The strategies for player 2 are defined analogously. We denote by Σ and Π the set of all strategies for player 1 and player 2, respectively.

Once a starting state $s \in S$ and strategies $\sigma \in \Sigma$ and $\pi \in \Pi$ for the two players are fixed, the outcome of the game is a random walk $\omega_s^{\sigma,\pi}$ for which the probabilities of events are uniquely defined, where an *event* $\mathcal{A} \subseteq \Omega$ is a measurable set of paths. Given strategies σ for player 1 and π for player 2, a play $\omega = \langle s_0, s_1, s_2, \ldots \rangle$ is *feasible* if for every $k \in \mathbb{N}$ the following three conditions hold: (1) if $s_k \in S_{\bigcirc}$, then $(s_k, s_{k+1}) \in E$; (2) if $s_k \in S_1$, then $\sigma(s_0, s_1, \ldots, s_k)(s_{k+1}) > 0$; and (3) if $s_k \in S_2$ then $\pi(s_0, s_1, \ldots, s_k)(s_{k+1}) > 0$. Given two strategies $\sigma \in \Sigma$ and $\pi \in \Pi$, and a state $s \in S$, we denote by $\mathrm{Outcome}(s, \sigma, \pi) \subseteq \Omega_s$ the set of feasible plays that start from s given the strategies σ and π. For a state $s \in S$ and an event $\mathcal{A} \subseteq \Omega$, we write $\Pr_s^{\sigma,\pi}(\mathcal{A})$ for the probability that a path belongs to $\mathcal{A}$ if the game starts from the state s and the players follow the strategies σ and π, respectively.

Strategies that do not use randomization are called pure. A player-1 strategy σ is *pure* if for all $\boldsymbol{w} \in S^*$ and $s \in S_1$, there is a state $t \in S$ such that $\sigma(\boldsymbol{w} \cdot s)(t) = 1$. A *memoryless* player-1 strategy does not depend on the history of the play but only on the current state; it can be represented as a function $\sigma\colon S_1 \to \mathcal{D}(S)$. A *pure memoryless strategy* is a strategy that is both pure and memoryless. A pure memoryless strategy for player 1 can be represented as a function $\sigma\colon S_1 \to S$. We denote by Σ^{PM} the set of pure memoryless strategies for player 1. The pure memoryless player-2 strategies Π^{PM} are defined analogously.

Given a pure memoryless strategy $\sigma \in \Sigma^{PM}$, let G_σ be the game graph obtained from G under the constraint that player 1 follows the strategy σ. The corresponding definition G_π for a player-2 strategy $\pi \in \Pi^{PM}$ is analogous, and we write $G_{\sigma,\pi}$ for the game graph obtained from G if both players follow the pure memoryless strategies σ and π, respectively. Observe that given a $2^1/_2$-player game graph G and a pure memoryless player-1 strategy σ, the result G_σ is a player-2 MDP. Similarly, for a player-1 MDP G and a pure memoryless player-1 strategy σ, the result G_σ is a Markov chain. Hence, if G is a $2^1/_2$-player game graph and the two players follow pure memoryless strategies σ and π, the result $G_{\sigma,\pi}$ is a Markov chain.

Objectives. We specify objectives for the players by providing a set of *winning* plays $\Phi \subseteq \Omega$ for each player. We say that a play ω *satisfies* the objective Φ if $\omega \in \Phi$. We study only zero-sum games, where the objectives of the two players are complementary; i.e., if player 1 has the objective Φ, then player 2 has the objective $\Omega \setminus \Phi$. We consider *ω-regular objectives* [11], specified as parity conditions. We also define the special case of reachability objectives.

- *Reachability objectives.* Given a set $T \subseteq S$ of "target" states, the reachability objective requires that some state of T be visited. The set of winning plays is $\mathrm{Reach}(T) = \{\, \omega = \langle s_0, s_1, s_2, \ldots \rangle \in \Omega \mid s_k \in T \text{ for some } k \geq 0 \,\}$.
- *Parity objectives.* For $c, d \in \mathbb{N}$, we write $[c..d] = \{\, c, c+1, \ldots, d \,\}$. Let p: $S \to [0..d]$ be a function that assigns a *priority* $p(s)$ to every state $s \in S$, where $d \in \mathbb{N}$. For a play $\omega = \langle s_0, s_1, \ldots \rangle \in \Omega$, we define $\mathrm{Inf}(\omega) = \{\, s \in S \mid s_k = s \text{ for infinitely many } k \,\}$ to be the set of states that occur infinitely often in ω. The *even-parity objective* is defined as $\mathrm{Parity}(p) = \{\, \omega \in \Omega \mid \min\big(p(\mathrm{Inf}(\omega))\big) \text{ is even } \}$, and the *odd-parity objective* as $\mathrm{coParity}(p) = \{\omega \in \Omega \mid \min\big(p(\mathrm{Inf}(\omega))\big) \text{ is odd } \}$.

Sure winning, almost-sure winning, and optimality. Given a player-1 objective Φ, a strategy $\sigma \in \Sigma$ is *sure winning* for player 1 from a state $s \in S$ if for every strategy $\pi \in \Pi$ for player 2, we have $\mathrm{Outcome}(s, \sigma, \pi) \subseteq \Phi$. The strategy σ is *almost-sure winning* for player 1 from the state s for the objective Φ if for every player-2 strategy π, we have $\Pr_s^{\sigma,\pi}(\Phi) = 1$. The sure and almost-sure winning strategies for player 2 are defined analogously. Given an objective Φ, the *sure winning set* $\langle\langle 1 \rangle\rangle_{sure}(\Phi)$ for player 1 is the set of states from which player 1 has a sure winning strategy. The *almost-sure winning set* $\langle\langle 1 \rangle\rangle_{almost}(\Phi)$ for player 1 is the set of states from which player 1 has an almost-sure winning strategy. The sure winning set $\langle\langle 2 \rangle\rangle_{sure}(\Omega \setminus \Phi)$ and the almost-sure winning set $\langle\langle 2 \rangle\rangle_{almost}(\Omega \setminus \Phi)$ for player 2 are defined analogously. It follows from the definitions that for all $2^1/_2$-player game graphs and all objectives Φ, we have $\langle\langle 1 \rangle\rangle_{sure}(\Phi) \subseteq \langle\langle 1 \rangle\rangle_{almost}(\Phi)$. A game is sure (resp. almost-sure) winning for player i if player i wins surely (resp. almost-surely) from every state in the game.

Given objectives $\Phi \subseteq \Omega$ for player 1 and $\Omega \setminus \Phi$ for player 2, we define the *value* functions $\langle\langle 1 \rangle\rangle_{val}$ and $\langle\langle 2 \rangle\rangle_{val}$ for the players 1 and 2, respectively, as the following functions from the state space S to the interval $[0, 1]$ of reals: for all states $s \in S$, let $\langle\langle 1 \rangle\rangle_{val}(\Phi)(s) = \sup_{\sigma \in \Sigma} \inf_{\pi \in \Pi} \Pr_s^{\sigma,\pi}(\Phi)$ and $\langle\langle 2 \rangle\rangle_{val}(\Omega \setminus \Phi)(s) =$

$\sup_{\pi\in\Pi}\inf_{\sigma\in\Sigma}\Pr_s^{\sigma,\pi}(\Omega\setminus\Phi)$. In other words, the value $\langle\langle 1\rangle\rangle_{val}(\Phi)(s)$ gives the maximal probability with which player 1 can achieve her objective Φ from state s, and analogously for player 2. The strategies that achieve the value are called optimal: a strategy σ for player 1 is *optimal* from the state s for the objective Φ if $\langle\langle 1\rangle\rangle_{val}(\Phi)(s) = \inf_{\pi\in\Pi}\Pr_s^{\sigma,\pi}(\Phi)$. The optimal strategies for player 2 are defined analogously.

Consider a family $\Sigma^{\mathcal{C}} \subseteq \Sigma$ of special strategies for player 1. We say that the family $\Sigma^{\mathcal{C}}$ *suffices* with respect to a player-1 objective Φ on a class $\mathcal{G}$ of game graphs for *sure winning* if for every game graph $G\in\mathcal{G}$ and state $s\in\langle\langle 1\rangle\rangle_{sure}(\Phi)$, there is a player-1 strategy $\sigma\in\Sigma^{\mathcal{C}}$ such that for every player-2 strategy $\pi\in\Pi$, we have $\mathrm{Outcome}(s,\sigma,\pi)\subseteq\Phi$. Similarly, the family $\Sigma^{\mathcal{C}}$ *suffices* with respect to the objective Φ on the class $\mathcal{G}$ of game graphs for *almost-sure winning* if for every game graph $G\in\mathcal{G}$ and state $s\in\langle\langle 1\rangle\rangle_{almost}(\Phi)$, there is a player-1 strategy $\sigma\in\Sigma^{\mathcal{C}}$ such that for every player-2 strategy $\pi\in\Pi$, we have $\Pr_s^{\sigma,\pi}(\Phi)=1$; and for *optimality*, if for every game graph $G\in\mathcal{G}$ and state $s\in S$, there is a player-1 strategy $\sigma\in\Sigma^{\mathcal{C}}$ such that $\langle\langle 1\rangle\rangle_{val}(\Phi)(s) = \inf_{\pi\in\Pi}\Pr_s^{\sigma,\pi}(\Phi)$. We now state the classical determinacy results for 2-player and $2^1\!/_2$-player parity games.

Theorem 1 (Qualitative determinacy). [7] *For all 2-player game graphs and parity objectives Φ, we have $\langle\langle 1\rangle\rangle_{sure}(\Phi) = S\setminus\langle\langle 2\rangle\rangle_{sure}(\Omega\setminus\Phi)$. Moreover, on 2-player game graphs, the family of pure memoryless strategies suffices for sure winning with respect to parity objectives.*

Theorem 2 (Quantitative determinacy). [4] *For all $2^1\!/_2$-player game graphs, all parity objectives Φ, and all states s, we have $\langle\langle 1\rangle\rangle_{val}(\Phi)(s)+\langle\langle 2\rangle\rangle_{val}(\Omega\setminus\Phi)(s)=1$. The family of pure memoryless strategies suffices for optimality with respect to parity objectives on $2^1\!/_2$-player game graphs.*

Since in $2^1\!/_2$-player games with parity objectives pure memoryless strategies suffice for optimality, in the sequel we consider only pure memoryless strategies.

3 Strategy Improvement Algorithm

The main result of this paper is a strategy improvement algorithm for $2^1\!/_2$-player games with parity objectives. Before presenting the algorithm, we recall a few key properties of $2^1\!/_2$-player parity games, which were proved in [2, 3].

Useful properties. We first present a reduction of $2^1\!/_2$-player parity games to 2-player parity games, preserving the ability of player 1 to win almost-surely.

Reduction. Given a $2^1\!/_2$-player game graph $G=((S,E),(S_1,S_2,S_\bigcirc),\delta)$ with a priority function p: $S\to[0..d]$, we construct a 2-player game graph $\overline{G}=((\overline{S},\overline{E}),(\overline{S}_1,\overline{S}_2),\overline{\delta})$ together with a priority function $\overline{p}:\overline{S}\to[0..d]$. The construction is specified as follows. For every nonprobabilistic state $s\in S_1\cup S_2$, there is a corresponding state $\overline{s}\in\overline{S}$ such that (1) $\overline{s}\in\overline{S}_1$ iff $s\in S_1$, and (2) $\overline{p}(\overline{s})=p(s)$, and (3) $(\overline{s},\overline{t})\in\overline{E}$ iff $(s,t)\in E$. From the state $\overline{s}$ with $\overline{p}(\overline{s})=p(s)$, the players play the following 3-step game in $\overline{G}$. First, in state $\overline{s}$ player 2 chooses a successor $(\widetilde{s},2k)$, for $k\in\{0,1,\ldots,j\}$, where $p(s)=2j$ or $p(s)=2j-1$. For every state

$(\widetilde{s}, 2k)$, we have $\overline{p}(\widetilde{s}, 2k) = p(s)$. For $k > 1$, in state $(\widetilde{s}, 2k)$ player 1 chooses from two successors: state $(\widehat{s}, 2k-1)$ with $\overline{p}(\widehat{s}, 2k-1) = 2k-1$, or state $(\widehat{s}, 2k)$ with $\overline{p}(\widehat{s}, 2k) = 2k$. The state $(\widetilde{s}, 0)$ has only one successor $(\widehat{s}, 0)$, with $\overline{p}(\widehat{s}, 0) = 0$. Finally, in each state $(\widehat{s}, k)$ the choice is between all states $\overline{t}$ such that $(s,t) \in E$, and it belongs to player 1 if k is odd, and to player 2 if k is even. We denote by $\mathrm{Tr}_{almost}(G)$ the 2-player game graph $\overline{G}$, as defined by this reduction. Also given a pure memoryless strategy $\overline{\sigma}$ for the 2-player game graph $\overline{G}$, a strategy $\mathrm{Tr}_{almost}(\overline{\sigma}) = \sigma$ for the $2^1\!/\!_2$-player game graph G is defined as follows: $\sigma(s) = t$ iff $\overline{\sigma}(\overline{s}) = \overline{t}$, for all $s \in S_1$. Similar definitions hold for player 2.

Lemma 1. [3] *Given a $2^1\!/\!_2$-player game graph G with the parity objective* $\mathrm{Parity}(p)$ *for player 1, let $\overline{U}_1$ and $\overline{U}_2$ be the sure winning sets for players 1 and 2, respectively, in the 2-player game graph $\overline{G} = \mathrm{Tr}_{almost}(G)$ with the modified parity objective* $\mathrm{Parity}(\overline{p})$. *Define the sets U_1 and U_2 in the original $2^1\!/\!_2$-player game graph G by $U_1 = \{ s \in S \mid \overline{s} \in \overline{U}_1 \}$ and $U_2 = \{ s \in S \mid \overline{s} \in \overline{U}_2 \}$. Then the following assertions hold: (a) $U_1 = \langle\!\langle 1 \rangle\!\rangle_{almost}(\mathrm{Parity}(p)) = S \setminus U_2$; and (b) if $\overline{\sigma}$ is a pure memoryless sure winning strategy for player 1 from $\overline{U}_1$ in $\overline{G}$, then* $\mathrm{Tr}_{almost}(\overline{\sigma})$ *is an almost-sure winning strategy for player 1 from U_1 in G.*

Subgames. A set $U \subseteq S$ of states is *δ-closed* if for every probabilistic state $u \in U \cap S_{\bigcirc}$, if $(u,t) \in E$, then $t \in U$. The set U is *δ-live* if for every nonprobabilistic state $s \in U \cap (S_1 \cup S_2)$, there is a state $t \in U$ such that $(s,t) \in E$. A δ-closed and δ-live subset U of S induces a *subgame graph* of G, denoted by $G \upharpoonright U$.

Boundary probabilistic states. Given a set U of states, let $\mathrm{BP}(U) = \{ s \in U \cap S_{\bigcirc} \mid \exists t \in E(s).\, t \notin U \}$ be the set of *boundary* probabilistic states, which have an edge out of U. Given a set U of states and a parity objective $\mathrm{Parity}(p)$ for player 1, we define a transformation $\mathrm{Tr}_{win1}(U)$ of U as follows: every state s in $\mathrm{BP}(U)$ is converted to an *absorbing* state (a state with a self-loop) and assigned the even priority $2 \cdot \lfloor \frac{d}{2} \rfloor$; thus, every state in $\mathrm{BP}(U)$ is changed to a sure winning state for player 1. Observe that if U is δ-live, then $\mathrm{Tr}_{win1}(G \upharpoonright U)$ is a game graph.

Value classes. Given a parity objective Φ, for every real $r \in \mathbb{R}$ the *value class* with value r, denoted $\mathrm{VC}(r) = \{ s \in S \mid \langle\!\langle 1 \rangle\!\rangle_{val}(\Phi)(s) = r \}$, is the set of states with value r for player 1. It follows easily that for every $r > 0$, the value class $\mathrm{VC}(r)$ is δ-live. The following lemma establishes a connection between value classes, the transformation Tr_{win1}, and the almost-sure winning states.

Lemma 2. [2] *For every real $r > 0$, for the value class* $\mathrm{VC}(r)$ *with parity objective* $\mathrm{Parity}(p)$ *for player 1, the game* $\mathrm{Tr}_{win1}(G \upharpoonright \mathrm{VC}(r))$ *is almost-sure winning for player 1.*

It follows from Lemma 1 and Lemma 2 that for every value class $\mathrm{VC}(r)$ with $r > 0$, the game $\mathrm{Tr}_{almost}(\mathrm{Tr}_{win1}(G \upharpoonright \mathrm{VC}(r)))$ is sure winning for player 1.

Strategy improvement algorithm. Given a strategy π and a set U of states, we denote by $\pi \upharpoonright U$ a strategy that for every state in U follows the strategy π.

Values and value classes given by strategies. Given a player-2 strategy π and a parity objective Φ for player 1, we denote the value of player 1 given the strategy

π as follows: $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) = \sup_{\sigma\in\Sigma^{PM}} \Pr^{\sigma,\pi}_s(\Phi)$. Similarly, we define the value classes given strategy π as $\mathrm{VC}^{\pi}(r) = \{ s \in S \mid \langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) = r \}$, for all $r \in \mathbb{R}$.

Witnesses for player 2. Given a $2^1/_2$-player game graph G and a parity objective Φ for player 1, a *witness* $w_2 = (\pi, \overline{\pi}_Q)$ for player 2 is specified as follows: (a) the strategy π is a player-2 strategy for the $2^1/_2$-player game graph G; and (b) for every value class $\mathrm{VC}^{\pi}(r)$, the strategy $\overline{\pi}_Q \restriction \mathrm{VC}^{\pi}(r)$ is a player-2 strategy for the 2-player game graph $\overline{G}_r = \mathrm{Tr}_{almost}(\mathrm{Tr}_{win1}(G \restriction \mathrm{VC}^{\pi}(r)))$. We require that $\pi = \mathrm{Tr}_{almost}(\overline{\pi}_Q)$. The witness $w_2 = (\pi, \overline{\pi}_Q)$ for player 2 is an *optimal* witness if π is an optimal strategy for player 2.

Ordering of witnesses. We define an ordering relation $\prec$ on witnesses as follows: given two witnesses $w_2 = (\pi, \overline{\pi}_Q)$ and $w'_2 = (\pi', \overline{\pi}'_Q)$ for player 2, let $w_2 \prec w'_2$ iff one of the following two conditions holds:

1. for all states s, we have $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) \geq \langle\langle 1\rangle\rangle^{\pi'}_{val}(\Phi)(s)$, and for some state s, we have $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) > \langle\langle 1\rangle\rangle^{\pi'}_{val}(\Phi)(s)$; or
2. for all states s, we have $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) = \langle\langle 1\rangle\rangle^{\pi'}_{val}(\Phi)(s)$, and in some value class $\mathrm{VC}^{\pi}(r) = \mathrm{VC}^{\pi'}(r)$, we have $(\overline{\pi}_Q \restriction \mathrm{VC}^{\pi}(r)) \prec_Q (\overline{\pi}'_Q \restriction \mathrm{VC}^{\pi}(r))$ in the 2-player parity game $\mathrm{Tr}_{almost}(\mathrm{Tr}_{win1}(G \restriction \mathrm{VC}^{\pi}(r)))$, where $\prec_Q$ denotes the ordering of strategies for a strategy improvement algorithm for 2-player parity games (e.g., as defined in [1, 12]).

Profitable switches. Given a witness $w_2 = (\pi, \overline{\pi}_Q)$ for player 2, we specify a procedure `ProfitableSwitch` to "improve" the witness according to the witness ordering $\prec$. The procedure is described in Algorithm 1. An informal description of the procedure is as follows: given a witness $w_2 = (\pi, \overline{\pi}_Q)$, the algorithm computes the values $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s)$ for all states. If there is a state $s \in S_2$ such that the strategy can be "value improved," i.e., there is a state $t \in E(s)$ with $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(t) < \langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s)$, then the witness is modified by setting $\pi(s)$ to t. This step is similar to the strategy improvement step of [6] and is achieved in Step 2.1 of `ProfitableSwitch`. Otherwise, in every value class $\mathrm{VC}^{\pi}(r)$, the strategy $\overline{\pi}_Q$ is "improved" for the game $\mathrm{Tr}_{almost}(\mathrm{Tr}_{win1}(G \restriction \mathrm{VC}^{\pi}(r)))$ with respect to the ordering $\prec_Q$ of strategies for 2-player parity games. This is achieved in Step 2.2.

Proposition 1. *Given a strategy π for player 2, for all states $s \in \mathrm{VC}^{\pi}(r) \cap S_1$ and $t \in E(s)$, we have $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(t) \leq r$, that is, $E(s) \subseteq \bigcup_{0\leq q\leq r} \mathrm{VC}^{\pi}(q)$.*

Proposition 2. *Given a strategy π for player 2, for all strategies σ for player 1, if there is a closed recurrent class C in the Markov chain $G_{\sigma,\pi}$ with $C \subseteq \mathrm{VC}^{\pi}(r)$ for $r < 1$, then $\min(p(C))$ is odd.*

Lemma 3. *Let $w_2 = (\pi, \overline{\pi}_Q)$ be an input to Algorithm 1, and let $w'_2 = (\pi', \overline{\pi}'_Q)$ be the corresponding output, that is, $w'_2 =$ `ProfitableSwitch`(G, w_2). If the set I in Step 2 of Algorithm 1 is nonempty, then (a) $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) \geq \langle\langle 1\rangle\rangle^{\pi'}_{val}(\Phi)(s)$ for all $s \in S$, and (b) $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) > \langle\langle 1\rangle\rangle^{\pi'}_{val}(\Phi)(s)$ for all $s \in I$.*

Proof. Consider a switch of the strategy of player 2 from π to π', as constructed in Step 2.1 of Algorithm 1. Consider a strategy σ for player 1 and a closed

Algorithm 1 `ProfitableSwitch`

Input: $2^1/_2$-player game G, parity objective Φ for pl. 1, witness $w_2 = (\pi, \overline{\pi}_Q)$ for pl. 2.
Output: a witness w_2' for player 2 such that either $w_2 = w_2'$ or $w_2 \prec w_2'$.
1. (Step 1.) Compute $\langle\langle 1 \rangle\rangle^{\pi}_{val}(\Phi)(s)$ for all states s.
2. (Step 2.) Consider the set $I = \{ s \in S_2 \mid \exists t \in E(s).\ \langle\langle 1 \rangle\rangle^{\pi}_{val}(\Phi)(s) > \langle\langle 1 \rangle\rangle^{\pi}_{val}(\Phi)(t) \}$.
   2.1 (Value improvement) **if** $I \neq \emptyset$ **then** set π' as follows:
      $\pi'(s) = \pi(s)$ for $s \in S_2 \setminus I$,
      $\pi'(s) = t$ for $s \in I$ and $t \in E(s)$ such that $\langle\langle 1 \rangle\rangle^{\pi}_{val}(\Phi)(s) > \langle\langle 1 \rangle\rangle^{\pi}_{val}(\Phi)(t)$;
      and set $\overline{\pi}_Q'$ to be an arbitrary strategy such that $\pi' = \mathrm{Tr}_{almost}(\overline{\pi}_Q')$.
   2.2 (Qualitative improvement) **else** for every value class $\mathrm{VC}^{\pi}(r)$:
      let $\overline{G}_r$ be the 2-player game graph $\mathrm{Tr}_{almost}(\mathrm{Tr}_{win1}(G \upharpoonright \mathrm{VC}^{\pi}(r)))$;
      set $(\overline{\pi}_Q' \upharpoonright \mathrm{VC}^{\pi}(r)) = \texttt{TwoPlSwitch}(\overline{G}_r,\ \overline{\pi}_Q \upharpoonright \mathrm{VC}^{\pi}(r))$ and $\pi' = \mathrm{Tr}_{almost}(\overline{\pi}_Q')$,
      where `TwoPlSwitch` is a strategy improvement step for 2-player parity games.
3. **return** $w_2' = (\pi', \overline{\pi}_Q')$.

---

recurrent class C in $G_{\sigma,\pi'}$ such that $C \subseteq \bigcup_{r<1} \mathrm{VC}^{\pi}(r)$. Let $z = \min\{ r < 1 \mid C \cap \mathrm{VC}^{\pi}(r) \neq \emptyset \}$, that is, $\mathrm{VC}^{\pi}(z)$ is the least value class with nonempty intersection with C. A state $s \in \mathrm{VC}^{\pi}(z) \cap C$ satisfies the following conditions:

1. If $s \in S_1$, then $\sigma(s) \in \mathrm{VC}^{\pi}(z)$. This follows since, by Proposition 1, we have $E(s) \subseteq \bigcup_{0 \le q \le z} \mathrm{VC}^{\pi}(q)$ and $C \cap \mathrm{VC}^{\pi}(q) = \emptyset$ for $q < z$.
2. If $s \in S_2$, then $\pi'(s) \in \mathrm{VC}^{\pi}(z)$. This follows since, by construction, we have $\pi'(s) \in \bigcup_{0 \le q \le z} \mathrm{VC}^{\pi}(q)$ and $C \cap \mathrm{VC}^{\pi}(q) = \emptyset$ for $q < z$. Also, since $s \in \mathrm{VC}^{\pi}(z)$ and $\pi'(s) \in \mathrm{VC}^{\pi}(z)$, it follows that $\pi'(s) = \pi(s)$.
3. If $s \in S_{\bigcirc}$, then $E(s) \subseteq \mathrm{VC}^{\pi}(z)$. This follows since for $s \in S_{\bigcirc}$, if $E(s) \subsetneq \mathrm{VC}^{\pi}(z)$, then $E(s) \cap (\bigcup_{0 \le q < z} \mathrm{VC}^{\pi}(q)) \neq \emptyset$. Since C is closed, and $C \cap \mathrm{VC}^{\pi}(q) = \emptyset$ for $q < z$, the claim follows.

It follows that $C \subseteq \mathrm{VC}^{\pi}(z)$ and for all states $s \in C \cap S_2$, we have $\pi'(s) = \pi(s)$. Hence by Proposition 2, $\min(p(C))$ is odd.

It follows that if player 2 switches to the strategy π', as constructed when Step 2.1 of Algorithm 1 is executed, then for all strategies σ for player 1 the following assertion holds: if there is a closed recurrent class $C \subseteq S \setminus \mathrm{VC}^{\pi}(1)$ in the Markov chain $G_{\sigma,\pi'}$, then C is winning for player 2, i.e., $\min(p(C))$ is odd. Hence, given strategy π', an optimal counter-strategy for player 1 maximizes the probability to reach $\mathrm{VC}^{\pi}(1)$. The desired result follows from arguments similar to $2^1/_2$-player games with reachability objectives [6], with $\mathrm{VC}^{\pi}(1)$ as the target for player 1, and the value improvement step (Step 2.1 of Algorithm 1). ∎

Lemma 4. *Let $w_2 = (\pi, \overline{\pi}_Q)$ be an input to Algorithm 1, and let $w_2' = (\pi', \overline{\pi}_Q')$ be the corresponding output, that is, $w_2' =$ `ProfitableSwitch`(G, w_2), such that $w_2 \neq w_2'$. If the set I in Step 2 of Algorithm 1 is empty, then (a) $\langle\langle 1 \rangle\rangle^{\pi}_{val}(\Phi)(s) \geq \langle\langle 1 \rangle\rangle^{\pi'}_{val}(\Phi)(s)$ for all $s \in S$, and (b) if $\langle\langle 1 \rangle\rangle^{\pi}_{val}(\Phi)(s) = \langle\langle 1 \rangle\rangle^{\pi'}_{val}(\Phi)(s)$ for all $s \in S$, then $(\overline{\pi}_Q \upharpoonright \mathrm{VC}^{\pi}(r)) \prec_Q (\overline{\pi}_Q' \upharpoonright \mathrm{VC}^{\pi}(r))$ for some value class $\mathrm{VC}^{\pi}(r)$.*

Proof. (sketch) It follows from Proposition 2 that for all strategies σ for player 1, if C is a closed recurrent class in $G_{\sigma,\pi}$ and $C \subseteq \mathrm{VC}^{\pi}(r)$ for $r < 1$, then $\min(p(C))$ is odd. Let π' be the strategy constructed from π in Step 2.2 of Algorithm 1. Since π' is obtained as qualitative improvement of π, it can be shown that if C is a closed recurrent class in $G_{\sigma,\pi'}$ and $C \subseteq \mathrm{VC}^{\pi}(r)$, then $\min(p(C))$ is odd. Arguments similar to Lemma 3 show that: for all strategies σ for player 1, if there is a closed recurrent class $C \subseteq S \setminus \mathrm{VC}^{\pi}(1)$ in $G_{\sigma,\pi'}$, then C is winning for player 2, that is, $\min(p(C))$ is odd. Since in strategy π' player 2 chooses every edge in the same value class as π, it can be shown that for all states s, we have $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) \geq \langle\langle 1\rangle\rangle^{\pi'}_{val}(\Phi)(s)$. If $\langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) = \langle\langle 1\rangle\rangle^{\pi'}_{val}(\Phi)(s)$ for all states s, then $\mathrm{VC}^{\pi}(r) = \mathrm{VC}^{\pi'}(r)$ for all r, that is, the value classes given π and π' coincide. Then, by the properties of the procedure `TwoPlSwitch` and since $w_2 \neq w_2'$, condition 2 of the lemma holds. ∎

Lemma 5. *For every $2^1/_2$-player game graph G and witness w_2 for player 2, if $w_2 \neq$ `ProfitableSwitch`(G, w_2), then $w_2 \prec$ `ProfitableSwitch`(G, w_2).*

Lemma 3 and Lemma 4 yield Lemma 5. The key argument to establish that if $w_2 =$ `ProfitableSwitch`(G, w_2), then w_2 is an optimal witness for player 2, goes as follows. Let w_2 be a player-2 witness such that $w_2 =$ `ProfitableSwitch`(G, w_2), and let $w_1 = (\sigma, \overline{\sigma}_Q)$ be the optimal counter-witness for player 1. Consider a value class $\mathrm{VC}^{\pi}(r)$ with $r > 0$, and the game graph $\overline{G}_r = \mathrm{Tr}_{almost}(\mathrm{Tr}_{win1}(G \upharpoonright \mathrm{VC}^{\pi}(r)))$. Since $\overline{\pi}_Q$ cannot be improved against $\overline{\sigma}_Q$ with respect to the ordering $\prec_Q$ in any value class, it follows that $\overline{\sigma}_Q$ is a sure winning strategy in $\overline{G}_r$. Hence it follows from Lemma 1 that σ is an almost-sure winning strategy for player 1 in $\mathrm{Tr}_{win1}(G \upharpoonright \mathrm{VC}^{\pi}(r))$, since $\sigma = \mathrm{Tr}_{almost}(\overline{\sigma}_Q)$. Consider any strategy π' for player 2 against σ, and consider the Markov chain $G_{\sigma,\pi'}$. Since σ is almost-sure winning in $\mathrm{Tr}_{win1}(G \upharpoonright \mathrm{VC}^{\pi}(r))$ for all $r > 0$, it follows that for any closed recurrent class C of $G_{\sigma,\pi'}$ such that $C \subseteq \bigcup_{r>0} \mathrm{VC}^{\pi}(r)$, the set C is winning for player 1 (i.e., $\min(p(C))$ is even). Moreover, since the strategy π cannot be "value improved," it follows from arguments similar to [6] for $2^1/_2$-player reachability games that for all player-2 strategies π' and all states $s \in \mathrm{VC}^{\pi}(r)$, we have $\mathrm{Pr}^{\sigma,\pi'}_s(\Phi) \geq r$. Hence $\langle\langle 1\rangle\rangle_{val}(\Phi)(s) \geq r$. Since σ is an optimal strategy against π, we have $r = \langle\langle 1\rangle\rangle^{\pi}_{val}(\Phi)(s) \geq \langle\langle 1\rangle\rangle_{val}(\Phi)(s)$ for all states $s \in \mathrm{VC}^{\pi}(r)$. This establishes the optimality of π.

Lemma 6. *For every $2^1/_2$-player game graph G and witness w_2 for player 2, if $w_2 =$ `ProfitableSwitch`(G, w_2), then w_2 is an optimal witness.*

A strategy improvement algorithm using the `ProfitableSwitch` procedure is described in Algorithm 2. The correctness of the algorithm follows from Lemma 6. Let $I_2(k)$ and $I_R(k)$ denote bounds on the number of iterations of strategy improvement algorithms for 2-player parity games and $2^1/_2$-player reachability games, respectively, for game graphs with k states. The number of iterations of Algorithm 2 between two value improvement steps can be bounded by $n \cdot I_2(n \cdot d)$, and hence the total number of iterations of Algorithm 2 can be bounded by $n \cdot I_2(n \cdot d) \cdot I_R(n \cdot d)$. Given an optimal strategy π for player 2,

Algorithm 2 StrategyImprovementAlgorithm

Input: a $2^1\!/_2$-player game graph G with parity objective Φ for player 1.
Output: a witness w_2^* for player 2.
1. Choose an arbitrary witness w_2 for player 2.
2. **while** $w_2 \neq \texttt{ProfitableSwitch}(G, w_2)$ **do** $w_2 = \texttt{ProfitableSwitch}(G, w_2)$.
3. **return** $w_2^* = w_2$.

the values for both the players can be computed in polytime by computing the values of the MDP G_π [4].

Theorem 3 (Correctness of Algorithm 2). *The output w_2^* of Algorithm 2 is an optimal witness for player 2.*

4 Randomized Subexponential Algorithm

We now present a randomized subexponential-time algorithm for $2^1\!/_2$-player parity games, by combining an algorithm of Björklund et al. [1] and the witness improvement procedure `ProfitableSwitch`.

Games and improving subgames. Given $l, m \in \mathbb{N}$, let $\mathcal{G}(l, m)$ be the class of $2^1\!/_2$-player game graphs with the set S_2 of player 2 states partitioned into two sets as follows: (a) $O_1 = \{ s \in S_2 \mid |E(s)| = 1 \}$, i.e., the set of states with out-degree 1; and (b) $O_2 = S_2 \setminus O_1$, with $O_2 \leq l$ and $\sum_{s \in O_2} |E(s)| \leq m$. There is no restriction for player 1. Given a game $G \in \mathcal{G}(l, m)$, a state $s \in O_2$, and an edge $e = (s, t)$, we define the subgame $\widetilde{G}_e$ by deleting all edges from s other than the edge e. Observe that $\widetilde{G}_e \in \mathcal{G}(l-1, m-|E(s)|)$, and hence also $\widetilde{G}_e \in \mathcal{G}(l, m)$. If $w_2 = (\pi, \overline{\pi}_Q)$ is a witness for player 2 in $G \in \mathcal{G}(l, m)$, then a subgame $\widetilde{G}$ is *w_2-improving* if some witness $w_2' = (\pi', \overline{\pi}_Q')$ in $\widetilde{G}$ satisfies $w_2 \prec w_2'$.

Informal description of Algorithm 3. The algorithm takes a $2^1\!/_2$-player parity game and an initial witness w_2^0, and proceeds in three steps. In Step 1, it constructs r pairs of w_2^0-improving subgames $\widetilde{G}$ and corresponding improved witnesses w_2 in $\widetilde{G}$. This is achieved by the procedure `ImprovingSubgames`. The parameter r will be chosen to obtain a suitable complexity analysis. In Step 2, the algorithm selects uniformly at random one of the improving subgames $\widetilde{G}$ with corresponding witness w_2, and recursively computes an optimal witness w_2^* in $\widetilde{G}$ from w_2 as the initial witness. If the witness w_2^* is optimal in the original game G, then the algorithm terminates and returns w_2^*. Otherwise it improves w_2^* by a call to `ProfitableSwitch`, and continues at Step 1 with the improved witness `ProfitableSwitch`(G, w_2^*) as the initial witness.

The procedure `ImprovingSubgames` constructs a sequence of game graphs $G^0, G^1, \ldots, G^{r-l}$ with $G^i \in \mathcal{G}(l, l+i)$ such that all $(l+i)$-subgames $\widetilde{G}_e^i$ of G^i are w_2^0-improving. The subgame G^{i+1} is constructed from G^i as follows: we compute

Algorithm 3 RandomizedAlgorithm ($2\,1/2$-player parity games)

Input: a $2\,1/2$-player game graph $G \in \mathcal{G}(l,m)$, a parity objective Parity(p) for pl. 1 and an initial witness w_2^0 for pl. 2.
Output : an optimal witness $w_2^* = (\pi^*, \overline{\pi}_Q^*)$ for player 2.
1. (Step 1) Collect a set I of r pairs $(\widetilde{G}, w_2)$ of subgames $\widetilde{G}$ of G, and corresponding witnesses w_2 in $\widetilde{G}$ such that $w_2^0 \prec w_2$.
(This is achieved by the procedure `ImprovingSubgames` below).
2. (Step 2) Select a pair $(\widetilde{G}, w_2)$ from I uniformly at random.
2.1 Find an optimal witness in $w_2^* \in \widetilde{G}$ by applying the algorithm recursively, with w_2 as the initial witness.
3. (Step 3) **if** w_2^* is an optimal witness in the original game G **then return** w_2^*.
else let $w_2 =$ `ProfitableSwitch`(G, w_2^*), and
goto Step 1 with G and w_2 as the initial witness.

procedure `ImprovingSubgames`
1. Construct sequence $G^0, G^1, \ldots, G^{r-l}$ of subgames with $G^i \in \mathcal{G}(l, l+i)$ as follows:
1.1 G^0 is the game where each edge is fixed according to w_2^0.
1.2 Let w_2^i be an optimal witness in G^i;
1.2.1 **if** w_2^i is an optimal witness in the original game G
then return w_2^i.
1.2.2 **else** let e be any target of `ProfitableSwitch`(G, w_2^i);
the subgame G^{i+1} is G^i with the edge e added.
2. **return** r subgames (fixing one of the r edges in G^{r-l}) and associated witnesses.

an optimal witness w_2^i in G^i, and if w_2^i is optimal in G, then we have discovered an optimal witness; otherwise we construct G^{i+1} by adding any *target* edge e of `ProfitableSwitch`(G, w_2^i) in G^i, i.e., e is an edge required in the witness `ProfitableSwitch`(G, w_2^i) that is not in the witness w_2^i.

The correctness of the algorithm can be seen as follows. Observe that every time Step 1 is executed, the initial witness is improved with respect to the ordering $\prec$ on witnesses. Since the number of witnesses is bounded, the termination of the algorithm is guaranteed. Step 3 of Algorithm 3 and Step 1.2.1 of procedure `ImprovingSubgames` ensure that on termination of the algorithm, the returned witness is optimal. Lemma 7 bounds the expected number of iterations of Algorithm 3. The analysis is similar to the results of [1].

Lemma 7. *Algorithm 3 computes an optimal witness. The expected number of iterations $T(\cdot,\cdot)$ of Algorithm 3 for a game $G \in \mathcal{G}(l,m)$ is bounded by the following recurrence:* $T(l,m) \leq \sum_{i=l}^{r} T(l,i) + T(l-1, m-2) + \frac{1}{r} \cdot \sum_{i=1}^{r} T(l, m-i) + 1$.

For a game graph G with $|S| = n$, we obtain a bound of n^2 for m. Applying a symmetric version of Algorithm 3 for player 1 if $|S_1| \leq |S_2|$, we can bound l by $\min\{\,|S_1|, |S_2|\,\}$. Using this fact and an analysis of Kalai for linear programming, Björklund et al. [1] showed that $m^{O\left(\sqrt{l/\log(l)}\right)} = 2^{O\left(\sqrt{l \cdot \log(l)}\right)}$ is a solution to the recurrence of Lemma 7, by choosing $r = \max\{\, l, \frac{m}{2}\,\}$, where $l = \min\{\,|S_1|, |S_2|\,\}$.

Lemma 8. *Procedure* `ProfitableSwitch` *can be computed in polynomial time.*

A call to `ProfitableSwitch` requires solving an MDP with parity objectives quantitatively (Step 1 of `ProfitableSwitch`; for a polynomial-time procedure, see [4]) and computing a profitable switch for 2-player parity games (Step 2.2 of `ProfitableSwitch`; for a polynomial-time procedure, see [1, 12]). Thus Lemma 8 follows. We obtain Theorem 4 as follows. Observe that the reduction Tr_{almost} of $2^1/_2$ player games to 2-player games causes a blow-up by a factor of d for states in $S_{\bigcirc}$. This fact, along with the solution for the recurrence of Lemma 7, using $l = d \cdot n_0 + \min\{n_1, n_2\}$ in the solution, yields that the expected number of iterations of Algorithm 3 is bounded by $2^{O\left(\sqrt{z \cdot \log(z)}\right)}$, where $z = d \cdot n_0 + \min\{n_1, n_2\}$. Each iteration of the algorithm requires a call to `ProfitableSwitch`. This analysis with Lemma 8 proves Theorem 4.

Theorem 4. *Given a $2^1/_2$-player game graph G with a priority function $p\colon S \to [0..d]$, the value $\langle\langle 1 \rangle\rangle_{val}(\mathrm{Parity}(p))(s)$ can be computed for all states $s \in S$ in time $2^{O\left(\sqrt{z \cdot \log(z)}\right)} \cdot O\big(poly(n)\big)$, where $n_1 = |S_1|$, $n_2 = |S_2|$, $n_0 = |S_{\bigcirc}|$, $z = (n_0 \cdot d + \min\{n_1, n_2\})$, and poly represents a polynomial function.*

Acknowledgments. This research was supported in part by the AFOSR MURI grant F49620-00-1-0327 and the NSF ITR grant CCR-0225610.

References

1. H. Bjorklund, S. Sandberg, and S. Vorobyov. A discrete subexponential algorithm for parity games. In *STACS*, pages 663–674. LNCS 2607, Springer, 2003.
2. K. Chatterjee, L. de Alfaro, and T.A. Henzinger. The complexity of stochastic Rabin and Streett games. In *ICALP*, pages 878–890. LNCS 3580, Springer, 2005.
3. K. Chatterjee, M. Jurdziński, and T. A. Henzinger. Simple stochastic parity games. In *CSL*, pages 100–113. LNCS 2803, Springer, 2003.
4. K. Chatterjee, M. Jurdziński, and T.A. Henzinger. Quantitative stochastic parity games. In *SODA*, pages 114–123. SIAM, 2004.
5. A. Condon. The complexity of stochastic games. *Information and Computation*, 96:203–224, 1992.
6. A. Condon. On algorithms for simple stochastic games. In *Advances in Computational Complexity Theory*, pages 51–73. American Mathematical Society, 1993.
7. E.A. Emerson and C. Jutla. The complexity of tree automata and logics of programs. In *FOCS*, pages 328–337. IEEE Computer Society Press, 1988.
8. A. Hoffman and R. Karp. On nonterminating stochastic games. *Management Science*, 12:359–370, 1966.
9. M. Jurdziński, M. Paterson, and U. Zwick. A deterministic subexponential algorithm for solving parity games. In *SODA (To appear)*, 2006.
10. W. Ludwig. A subexponential randomized algorithm for the simple stochastic game problem. *Information and Computation*, 117:151–155, 1995.
11. W. Thomas. Languages, automata, and logic. In *Handbook of Formal Languages*, volume 3, Beyond Words, chapter 7, pages 389–455. Springer, 1997.
12. J. Vöge and M. Jurdziński. A discrete strategy improvement algorithm for solving parity games. In *CAV*, pages 202–215. LNCS 1855, Springer, 2000.